

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
24 April 2003 (24.04.2003)

PCT

(10) International Publication Number
WO 03/034242 A1

(51) International Patent Classification⁷: **G06F 13/14**,
12/00, H04B 7/212, H04L 12/56

(21) International Application Number: PCT/US02/05439

(22) International Filing Date: 21 February 2002 (21.02.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/977,602 12 October 2001 (12.10.2001) US

(71) Applicant: **SONICS, INC.** [US/US]; 2440 West El
Camino Real, Suite #620, Mountain View, CA 94040
(US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN,
YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR,
GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent
(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR,
NE, SN, TD, TG).

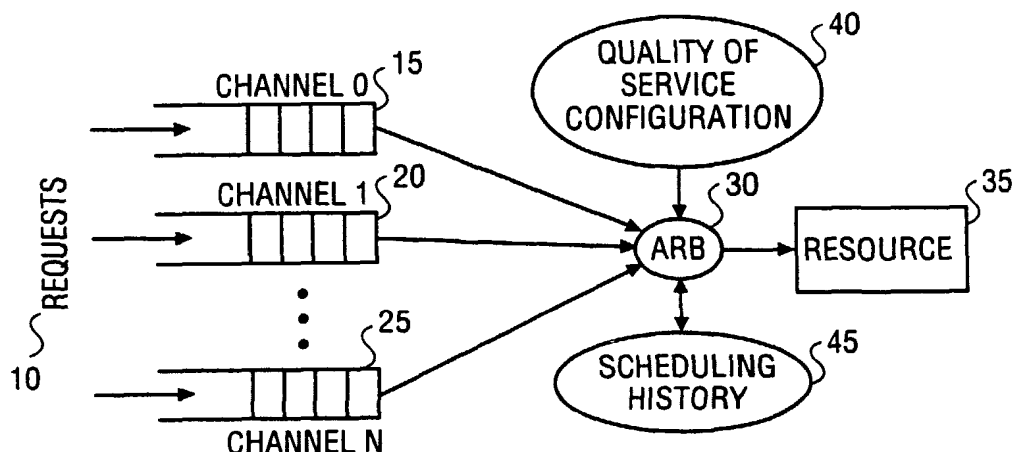
(72) Inventor: **WEBER, Wolf-Dietrich**; 5851 Meander Drive,
San Jose, CA 95120 (US).

Published:
— with international search report

(74) Agents: **SOBRINO, Maria, McCormack** et al.; Blakely,
Sokoloff, Taylor & Zafman, 7th Floor, 12400 Wilshire
Boulevard, Los Angeles, CA 90025-1026 (US).

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND APPARATUS FOR SCHEDULING A RESOURCE TO MEET QUALITY-OF-SERVICE RESTRICTIONS



(57) Abstract: The present invention is directed to a method and apparatus for scheduling a resource (35) to meet quality of service guarantees. In one embodiment of three levels of priority, if a channel of a first priority level (15) is within its bandwidth allocation, then a request is issued from that channel. If there are no requests in channels at the first priority level that are within the allocation, requests from channels at the second priority level (20) that are within their bandwidth allocation are chosen. If there are no requests of this type, requests from channels at the third priority level (25) or requests from channels at the first and second levels that are outside of their bandwidth allocation are issued. The system may be implemented using rate-based scheduling.



WO 03/034242 A1

METHOD AND APPARATUS FOR SCHEDULING A RESOURCE TO MEET QUALITY-OF-SERVICE RESTRICTIONS

FIELD OF THE INVENTION

[0001] The field of the invention relates to a system where access to a resource is scheduled to provide a particular quality-of-service to two or more requestors competing for access to that resource.

BACKGROUND

[0002] In computer systems it is common that a given resource (such as a system bus, a memory bank, etc.) is shared between several competing requesting devices or processes ("requestors") that would like to make use of the resource. Access to that resource therefore has to be arbitrated, in order to determine which requestor can access the resource when there are concurrent and conflicting requests to the resource. It is desirable to be able to specify different quality-of-service (QOS) guarantees for different requestors in order for the system to operate properly. Examples of QOS guarantees include data bandwidth and latency. For example, it may be desirable to allow a processor to have very high-priority and therefore low-latency access to a memory system. Another example is that one might want a video system to have a certain reserved bandwidth on a system bus so that the video screen can be updated as required at a fixed frame rate.

[0003] Existing arbitration schemes that aim to provide QOS guarantees include fixed-priority arbitration and time division multiplexing. In fixed-priority arbitration each requestor is assigned a fixed priority and requestors are serviced in priority order. In time division multiplexing, each requestor is pre-allocated a certain set of fixed access periods during which it can access the resource. While these arbitration schemes have their value in certain systems, they fall short of providing QOS guarantees when there is a mix of requestors with different QOS requirements and perhaps unpredictable request arrival times. For example, it is not possible to give any kind of bandwidth guarantee to multiple different requestors if fixed-priority arbitration is used unless the exact request pattern of each initiator is known a priori. Time division multiplexing is inefficient when the arrival times of requests are not deterministic, or when the requests require differing

amounts of service time from the resource depend on the type of request or the recent history of other requests.

[0004] What is desired is a resource scheduling scheme that can provide different QOS guarantees to different requestors and further can efficiently handle non-deterministic arrival and service times.

SUMMARY OF THE INVENTION

[0005] The present invention is directed to a method and apparatus for scheduling a resource to meet quality of service guarantees. In one embodiment of three levels of priority, if a channel of a first priority level is within its bandwidth allocation, then a request is issued from that channel. If there are no requests in channels at the first priority level that are within the allocation, requests from channels at the second priority level that are within their bandwidth allocation are chosen. If there are no requests of this type, requests from channels at the third priority level or requests from channels at the first and second levels that are outside of their bandwidth allocation are issued. The system may be implemented using rate-based scheduling.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The objects, features and advantages of the invention will be apparent from the following detailed description in which:

[0007] **Figure 1** is a simplified diagram of one embodiment of an arbitration system that operates in accordance with the technology of the present invention.

[0008] **Figure 2** illustrates one embodiment of priority order.

[0009] **Figure 3** is a simplified flow diagram illustrating one embodiment of an arbiter.

[0010] **Figure 4** illustrates an embodiment of rate-based scheduling.

[0011] **Figure 5** is a simplified diagram illustrating one embodiment of rate-based scheduling in accordance with the teachings of the present invention.

DETAILED DESCRIPTION

[0012] **Figure 1** shows one embodiment of an arbitration system. Requests 10 arrive from different requesting devices or processes and are stored in a channel, e.g., channels 15, 20, 25 that are contemplated to be logically or physically implemented. In this embodiment, each channel accommodates requests from one requestor. Thus, requests

from device A are stored in channel 0 (15), requests from device B are stored in channel 1 (20), etc. In the present embodiment, it is assumed here that requests within each channel are serviced in the order they are received in each channel, but this is not a necessary requirement for the invention described herein.

[0013] The arbitration unit 30 is responsible for scheduling access by each channel to the resource 35. A resource can be a variety of different apparatuses or processes, including memory and the like. The arbitration unit 30 is configured with the desired quality-of-service (QOS) guarantees for each of the channels using the QOS configuration unit 40. QOS may include a variety of criteria including one or more minimum, maximum or ranges of performance criteria for a particular dataflow, process or device. The unit 30 also keeps track of recent scheduling decisions using the scheduling history unit 45. Although the QOS unit 40 and scheduling history unit 45 are illustrated as separate units, it is readily apparent that the functionality of one or both of the quality of service configuration unit 40 and scheduling history unit 45 can be configured to be part of the arbitration unit 30 or joined into a single unit coupled to the arbitration unit 30. Further, it is contemplated that the one or more of the units 30, 40, 45 may be physically part of the resource 35.

[0014] If more than one channel 15, 20, 25 has a request waiting for service, the arbitration unit 30 selects the channel that can proceed to service using the scheduling history and desired QOS information retrieved respectively from the scheduling history unit 45 and quality of service configuration unit 40. The next request of the selected channel proceeds to access the resource and exits the system.

[0015] In one embodiment, the arbiter 30 uses the scheduling history 45 to determine if certain QOS guarantees can be met. For example, it is possible that the amount of time needed to access the resource depends on the relative timing of access of the resource or the type of request. For example, when accessing a bi-directional system bus, it may take longer for a write request to access the bus if it has recently been accessed with a read request, because the bus direction may need to be turned around first. This information may be determined from the scheduling history 45 and in turn affects the scheduling history 45. As noted above, the amount of time needed to access the resource may also depend on the type of request. For example, when accessing a dynamic

random access memory (DRAM) memory system, a request to an open DRAM page might take much less time than a request to a closed DRAM page.

[0016] In one embodiment, the different QOS modes used for scheduling may include priority service, allocated-bandwidth service and best-effort service. Each channel is assigned one QOS mode. For a channel to receive priority service or allocated-bandwidth service, it must be given a bandwidth allocation. Priority service provides bandwidth guarantees up to the allocated bandwidth, plus minimum latency service. Allocated-bandwidth service provides only bandwidth guarantees up to the allocated bandwidth. Best-effort service provides no QOS guarantees, and may actually receive no service at all. Additional quality-of-service modes are possible.

[0017] Further arbitration, e.g., using the scheduling history, may be used to determine selection of one of a plurality of pending requests at the same level of QOS. For example, if the scheduling history indicates the resource, e.g. a bus, has been operating in one direction, the arbiter may grant priority to requests arguing operation in that same direction.

[0018] In the present embodiment, each of the channels 15, 20, 25 are allocated one QOS mode and this information is placed in the QOS configuration unit 40. In order to allocate bandwidth to different channels, it is important to calculate the overall bandwidth available in the resource that is being accessed. This total bandwidth may be relatively easy to calculate, or it may depend on the request stream itself and must therefore be estimated using a particular expected request stream for a particular system. For example, when estimating the total available bandwidth of a DRAM system, it may be necessary to estimate the expected fraction of page hits and page misses in the request stream. When it is not possible to make a precise estimate of the total available bandwidth, a conservative estimate should be made that can be achieved under all or substantially all foreseeable conditions in the given system.

[0019] In the present embodiment, priority channels and allocated-bandwidth channels are all allocated a certain amount of bandwidth. In order to meet the QOS guarantees under all conditions, no more than the total available bandwidth of the resource should be allocated to these channels.

[0020] Channels using a priority QOS mode or allocated-bandwidth mode receive a higher QOS than channels that use a best-effort mode, but only while the channels

continue to operate within their respective bandwidth allocation. It is possible for a channel to request more than its allocated bandwidth, but in this case the QOS may degrade to below or equal best-effort channels.

[0021] **Figure 2** illustrates one example of the priority order in which channels get access to a resource. The top level 205 is reserved for priority channels that are within their bandwidth allocation. If there are any channels with requests in this category, they are provided access (serviced) as soon as possible, thus achieving low-latency access to the resource. The next lower level 210 is for allocated-bandwidth channels that are within their bandwidth allocation. Thus, if there are no eligible priority requests, allocated-bandwidth requests are serviced. Best-effort channels and priority or allocated-bandwidth channels that are outside of their allocated bandwidth are serviced with the lowest priority. These two groups can either be combined or serviced as two separate priorities 215, 220 as shown in **Figure 2**.

[0022] Using this scheduling method, allocated bandwidth and priority channels are substantially guaranteed to receive their allocated bandwidth. Amongst the two, priority channels are serviced with a higher priority, so these channels experience a lower access latency to the resource. If and when the priority and allocated-bandwidth channels are not using up the total available bandwidth of the resource, best-effort and other channels that are outside of their allocation can make use of the resource, thus ensuring that the resource does not sit idle while there are requests waiting.

[0023] **Figure 3** illustrates one embodiment of an arbitration process. For purposes of discussion, it is assumed that one channel is assigned to each level. However, it is contemplated that multiple channels can be operative at the same level. At step 300, if there are pending requests to be service by a particular resource, and resource bandwidth is available, step 305, a first level, for example, the level with the highest priority of service is examined, step 310. At step 315, if any requests from channels at the first level are within their bandwidth allocation, a request is issued from one channel of the first level, step 320. If requests from channels at the first level are not within allocation, requests from channels at the next level, which in one embodiment are channels at a level of a next-lower priority, are examined, step 325. If the channels at the level are within allocation, or alternately are not assigned an allocation bandwidth, step 330, a request is

issued from a channel of that level, step 335. This process continues, step 340, for each level of channels.

[0024] Utilizing the scheduling system of the present invention, the system can determine whether a given channel is operating within or beyond its allocated bandwidth. One way to implement this mechanism is described by the embodiment of **Figure 4**. As shown in **Figure 4**, time is divided into equal-sized scheduling periods, for example, periods 410, 420, 430. Bandwidth is allocated on the basis of a fixed number of requests per scheduling period. The scheduling unit decides to schedule those requests at any suitable time during the scheduling period. For purposes of discussion, this scheduling method is referred to as "rate-based" scheduling. By allowing the scheduling unit to schedule requests at any time during the scheduling period it does not rely on a known request arrival time. Furthermore, the scheduling unit is able to schedule requests so as to maximize the efficiency of the resource. In the example of **Figure 4**, four requests 412, 414, 416, 418 are always being scheduled in each scheduling period, but the exact time that each request is processed by the resource varies from scheduling period, e.g., period 420, to scheduling period, e.g., period 430.

[0025] One embodiment of rate-based scheduling is shown in **Figure 5**. The advantage of this embodiment is that it requires very little state—just two counters per channel. The mechanism shown is for one channel and is sufficient to determine whether the given channel is above or below its bandwidth allocation. Multiple versions of this can be implemented to support multiple channels. In addition, it is contemplated that the functionality described can be implemented a variety of ways.

[0026] A rate counter 510 is incremented at some small periodic interval, such as once every cycle. In the present embodiment, the rate counter 510 is configured to have a maximum value that is based on the allocated bandwidth of the channel. For example, if the bandwidth allocated to a particular channel is ten requests during each 100-cycle scheduling period, then the rate counter would be set up with a maximum value of 10.

[0027] Once the rate counter 510 reaches its maximum value, it causes the allocation counter 515 to be incremented, thus signaling that there is one more request "credit" available for that channel. The rate counter 510 resets and begins counting again. In one embodiment the rate counter is implemented as a simple register or location in memory with associated logic to test the value of the counter. Alternately, the overflow bit of the

counter may be used to increment the allocation counter 515 and reset the rate counter 510.

[0028] Each time a request is sent from that channel to the resource, the allocation counter is decremented, thus removing a credit. As long as the allocation count is positive, the channel is operating within its bandwidth allocation.

[0029] In one embodiment, the allocation count does not go beyond the number of requests allocated per scheduling period (either positive or negative). Saturation logic is included to insure that the allocation count does not exceed a specified saturation value. For example, a register or memory corresponding to counter 515 may include control logic that would not change the value in the counter beyond positive saturation value or negative saturation value. This enables the bandwidth use history to fade with time, allowing a channel to use more than its allocation during certain periods when bandwidth is available, while still maintaining the bandwidth guarantees when bandwidth becomes tight again.

[0030] While the above-described scheduling method is suitable for all kinds of systems that have multiple requestors competing for a shared resource, it is particularly well suited for shared dynamic random access memory (DRAM) memory systems. DRAM systems are especially difficult to schedule, because the service time of each request depends on the request type (e.g., read or write, burst size, etc.) and the request history which determines whether a particular request hits a page that is open or whether a page must first be opened before it can be accessed. Given a conservative estimate of the bandwidth that can be achieved with a certain set of request streams from different initiating (requestor) devices or processes, the described scheduling method can guarantee different QOS to different requestors while achieving a very high DRAM efficiency.

[0031] The invention has been described in conjunction with the preferred embodiment. It is evident that numerous alternatives, modifications, variations and uses will be apparent to those skilled in the art in light of the foregoing description.

CLAIMS**What is claimed is:**

1. A method comprising:
determining if a resource is available to service requests;
if resource bandwidth is available;
if at least one channel of a first level comprises at least one request and is within allocation, issuing at least one request from at the least one channel of the first level;
if the channel of the first level does not comprise at least one of at least one request and exceeds allocation, and if a channel of a second level comprises at least one request, issuing a request from the channel of the second level.
2. The method as set forth in claim 1, wherein if the channel of the first level does not comprise at least one of at least one request and exceeds allocation, and if the channel of the second level does not comprise at least one of at least one request and exceeds allocation, and a channel of a third level comprises at least one request, issuing a request from the channel of the third level.
3. The method as set forth in claim 1, wherein if the channel of the first level comprises at least one request and exceeds allocation, issuing requests from the channel of the second level until the channel of the first level is within allocation, wherein requests of channels of the first level are given priority when operating within allocation.
4. The method as set forth in claim 1, wherein if the channel of the first level exceeds allocation, reassigning the channel to a third level lower in priority to the first level and second level such that requests of the channel of the first level are issued after requests of the channel of the second level.
5. The method as set forth in claim 4, wherein if the channel originally of the first level and currently at the third level is operating within allocation, reassigning the channel to the first level.

6. The method as set forth in claim 1, further comprising referencing at least one scheduling history of requests, said issuing also based upon the at least one scheduling history.

7. The method as set forth in claim 1, wherein the levels are representative of Quality of Service guarantees (QOS).

8. The method as set forth in claim 1, wherein the levels are representative of channels characterized from the group consisting of priority channels, allocated bandwidth channels and best effort channels.

9. The method as set forth in claim 1, wherein at least one channel comprises multiple requests, wherein ordering of issuance of requests within a channel is substantially independent of the issuance of requests from the channel to the resource.

10. A method comprising:
determining if a resource is available to service requests;
if resource bandwidth is available;
if at least one channel of a first level comprises at least one request and is within allocation, issuing at least one request from at the least one channel of the first level;
if the channel of the first level does not comprise at least one of at least one request and exceeds allocation, and if a channel of a second level, lower in priority than the first level, comprises at least one request, issuing a request from the channel of the second level; and
if the channel of the first level does not comprise at least one of at least one request and exceeds allocation, and if the channel of a second level does not comprise at least one request, issuing a request from a channel of the first level which exceeds allocation.

11. The method as set forth in claim 10, wherein if the channel of the first level does not comprise at least one of at least one request and exceeds allocation, and if the channel of the second level does not comprise at least one of at least one request and exceeds allocation, and a channel of a third level comprises at least one request, the third level lower in priority than the first level and second level issuing a request from the channel of the third level.

12. The method as set forth in claim 10, wherein if the channel of the first level comprises at least one request and exceeds allocation, issuing requests from the channel of the second level until the channel of the first level is within allocation, wherein requests of channels of the first level are given priority when operating within allocation.

13. The method as set forth in claim 10, wherein if the channel of the first level exceeds allocation, reassigning the channel to a third level lower in priority to the first level and second level such that requests of the channel of the first level are issued after requests of the channel of the second level.

14. The method as set forth in claim 13, wherein if the channel originally of the first level and currently at the third level is operating within allocation, reassigning the channel to the first level.

15. The method as set forth in claim 11, wherein channels of the first level comprise priority channel, channels of the second level comprise allocated bandwidth channels and channels of the third level comprise best effort channels.

16. An arbiter coupled to receive requests from channels of a plurality of levels, each level comprising at least one channel configured to control conveyance of requests for a resource to process, said arbiter configured to determine if the resource is available to service requests, and if resource bandwidth is available and if at least one channel of a first level comprises at least one request and is within allocation, said arbiter controlling issuance of at least one request from at the least one channel of the first level;

and if the channel of the first level does not comprise at least one of at least one request and exceeds allocation, and if a channel of a second level comprises at least one request, said arbiter controlling issuance of at least one request from the channel of the second level.

17. The arbiter as set forth in claim 16, wherein each level represents a priority of Quality of Service (QOS) guarantees.

18. The arbiter as set forth in claim 16, further comprising a scheduling history and identifications of quality of service (QOS) guarantees for the channels, said arbiter further controlling issuance of at least one request based upon the scheduling history and QOS guarantees.

19. The arbiter as set forth in claim 16, wherein the scheduling history comprises bandwidth usage, said arbiter further comprising linked counters to keep track of a channel's bandwidth usage history.

20. The arbiter as set forth in claim 16, wherein the scheduling history comprises bandwidth usage, said arbiter further comprising saturation logic such that the history of bandwidth usage of a channel is within a range specified by a saturation usage value.

21. The arbiter as set forth in claim 16, wherein the linked counters comprise a rate counter and allocation counter, the rate counter incremented every unit cycle, the allocation counter incremented every determined number of rate count and decremented each time a request is received from the channel.

22. The arbiter as set forth in claim 16, wherein the arbiter is further configured, if the channel of the first level does not comprise at least one of at least one request and exceeds allocation, and if the channel of the second level does not comprise at least one of at least one request and exceeds allocation, and a channel of a third level comprises at least one request, to issue a request from the channel of the third level.

23. The arbiter as set forth in claim 16, wherein the arbiter arbitrates requests to a dynamic random access memory (DRAM).

24. A system comprising:
at least one priority channel comprising at least one request for service from a resource, said priority channel provided a first determined allocation;
at least one allocated-bandwidth channel comprising at least one request for service from the resource, said allocated-bandwidth channel provided a second determined allocation; and
an arbiter comprising logic configured to determine ordering of requests using a rate-based scheduling process.

25. The system as set forth in claim 24, wherein the resource is a dynamic random access memory (DRAM).

26. The system as set forth in claim 24, where the rate-based scheduling process comprises incrementing a rate count every unit cycle, incrementing a first allocation count every determined number of rate count and decrementing the first allocation count each time a priority channel request is received, wherein when the first allocation count is below a specified threshold, requests from the allocated-bandwidth channel are processed before requests from the priority channel.

27. The system as set forth in claim 26, wherein the threshold corresponds to a saturation value, the first allocation count saturating at least one of the positive saturation value and negative saturation value.

28. The system as set forth in claim 24, wherein the arbiter comprises a rate counter and allocation counter, the rate counter incremented every unit cycle, the allocation counter incremented every determined number of rate counts and decremented each time a request is received from a channel.

29. The system as set forth in claim 28, wherein the allocation counter saturates at a saturation value such that the allocation counter value is within a range defined by positive saturation value and negative saturation value.

1/3

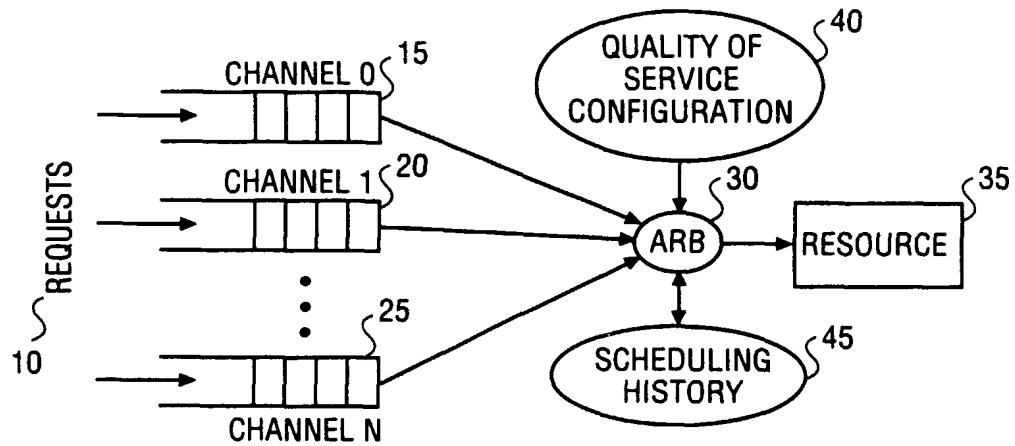


FIG. 1




FIG. 2

2/3

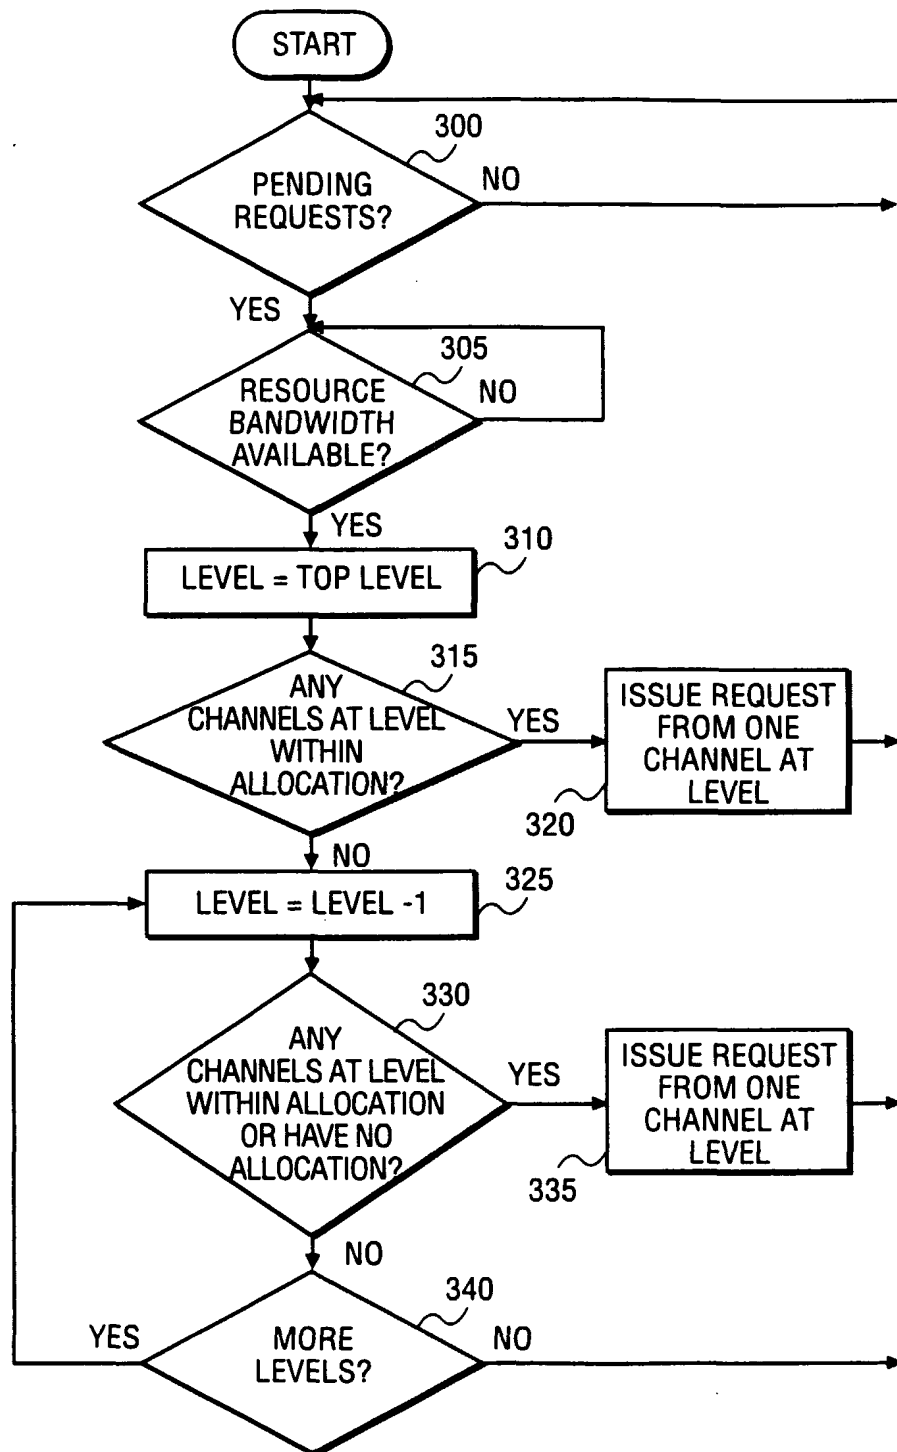


FIG. 3

3/3




FIG. 4

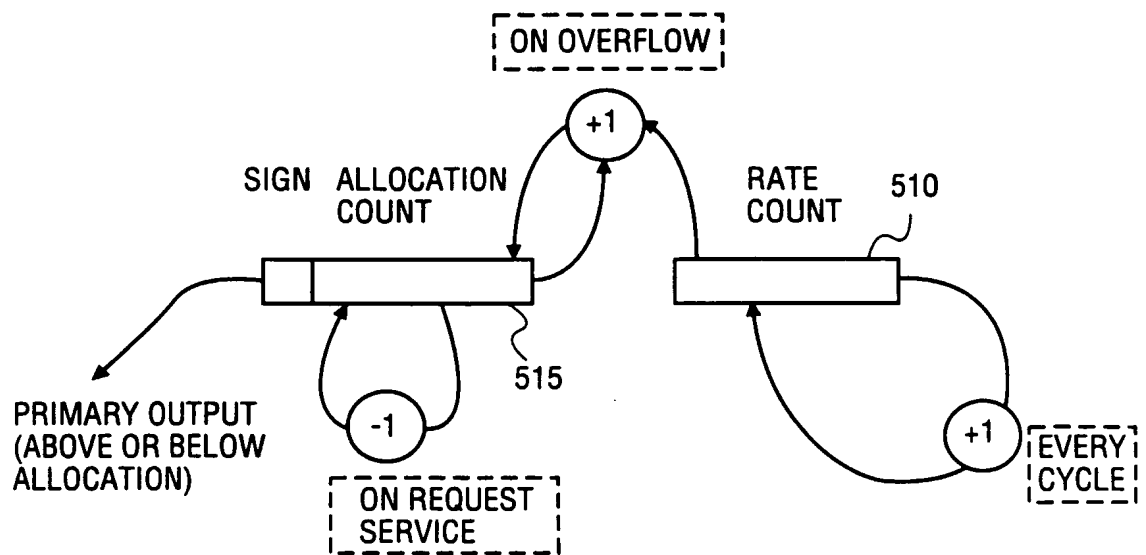


FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US02/05139

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 13/14, 12/00; H04B 7/212; H04L 12/56
US CL : 710/244, 117, 124, 45; 711/150, 151; 370/395.11, 395.12, 444
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 710/240-244, 107-125, 36-45; 711/147-151; 370/395.1, 395.11, 395.12, 443, 444

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
STN

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,218,456 A (STEGBAUER et al) 08 June 1993, abstract, col. 5, line 39 - col. 6, line 54.	1-5, 9-14, 16, 22, 24
X	US 5,996,037 A (EMNETT) 30 November 1999, abstract, col. 4, line 20 - col. 5, line 67.	1-5, 9-14, 16, 22, 24
A, E	US 6,363,445 B1 (JEDDELOH) 26 March 2002, abstract.	1-29
A	US 6,092,137 A (HUANG et al) 18 July 2000, abstract.	1-29
A	US 5,809,538 A (POLLMAN et al) 15 September 1998, abstract.	1-29

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	"I"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&"	document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means		
"P" document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search
29 MAY 2002

Date of mailing of the international search report
26 JUN 2002

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

SUMATI LEFKOWITZ *[Signature]*

Telephone No. (703) 308-7790